REMARKS

Claims 1-43 are presently pending. Claims 1-25 and 30-43 are rejected, while claims 26-29 are objected to. The Applicant requests further consideration and re-examination in view of the amendments above and remarks set forth below.

Rejections under 35 U.S.C. § 103(a):

Claims 1, 6-11 and 35 were rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,230,151 issued to Agrawal et al. (hereinafter "Agrawal et al.") in view of U.S. Patent No. 6,411,222 issued to Clark et al. (hereinafter "Clark et al.").

The Applicant respectfully traverses the rejection on the grounds that Agrawal et al. and Clark et al. cannot be properly combined. The Applicant also traverses on the grounds that, even if these references could be properly combined, such a combination does not suggest or disclose the Applicant's claimed invention.

Agrawal et al. relates to classification of data records for data mining purposes. Agrawal et al. at col. 1, lines 6-9 and 23-24. More particularly, Agrawal et al. discloses a method and apparatus for generating a decision tree classifier in a parallel multiprocessor system and in a shared memory multiprocessor system. Abstract of Agrawal. The decision tree classifier recursively partitions a set of example records, i.e. a training set, until each partition consists entirely or predominantly of examples from the same class. Agrawal et al. at col. 1, lines 40-44. In an example training set of data records, each record represents a car insurance applicant and includes three attributes: Age, Car Type and Risk Level. Agrawal et al. at col. 1, lines 52-54. A decision tree is constructed using this training set which can then be used to screen future applicants by classifying them as having a high or low Risk Level. Agrawal et al., at col. 1, lines 54-65.

The decision tree of Agrawal et al. contains tree-structured nodes in which each node is either a leaf, indicating a class, or a decision node, specifying some test on one or more data record attributes. Agrawal et al. at col. 7, lines 11-16. A decision tree classifier is usually built in two phases: a growth phase and a prune phase. Agrawal et al., at col. 7, lines 21-22. For the growth phase, given a set of training records, there are two cases to consider. Agrawal et al., at col. 7, lines 23-24. In the first case, if the records in the set are all of the same class, then the set becomes

a leaf. Agrawal et al., at col. 7, lines 24-26. Alternatively, in the second case, if the set contains records from more than one class, the set is partitioned into smaller sets based on an attribute test. Agrawal et al., at col. 7, lines 27-30. This process is performed recursively for each partitioned set of records until each partition consists entirely or predominantly of examples from the same class. Agrawal et al. at col. 1, lines 16-18 and 31-32.

Clark et al. discloses a generic problem modeler for examining an existing user information resource and transform the relevant information from the resource into data that can be stored in a solver database which is directly accessible by a problem solver. Clark et al., at col. 2, lines 49-54. This is intended to avoid the requirement that an application developer manually hard-code functions for the transformation of the data from the data source into properly formatted parameters to be interpreted by the problem solver. Clark et al., at col. 2, lines 44-48. Thus, an object of Clark et al. is to provide a generic problem modeler that requires less user skill than prior art systems and to provide methods for creating models for resource optimization problems. Clark et al., at col. 2, lines 55-58.

The problem modeler of Clark et al. determines metrics and metric relationships for each object of the information resource or may use defaults for those objects that it does not sufficiently understand. Clark et al., at col. 5, lines 28-32. A graphical user interface allows a non-expert developer to configure the system. Clark et al., at col. 5, lines 39-42. The user usually parses or breaks down the problem into smaller and discrete components and describes it in high-level terms. Clark et al., at col. 5, lines 42-44. The problem modeler uses information provided by the user to select and extract data from information resource and transform it to create optimization matrices, algorithms and constraints to be used by the problem solver for solving the user's optimization problem. Clark et al., at col. 5, lines 44-49. Clark et al. explain that by taking advantage of an existing database of related objects with defined relationship to each other for the information resource, much of the highly technical analysis performed by in the prior art by an application developer and a specific domain expert is circumvented. Clark at al., at col. 5, lines 59-67. Clark et al. provide an example embodiment of a resource optimization system for a telecommunications network design and assign application. Clark et al. at col. 7, line 14 to col. 9, line 64.

To establish a prima facie case of obviousness, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one skilled in the art to modify a reference or to combine reference teachings. Manual of Patent Examining Procedure, Section 2142 (Rev. 2, May 2004). The Examiner stated that Agrawal et al. does not explicitly teach a solver process, but that Clark et al. teaches a solver process and that it would have been obvious to combine the teaching of Clark et al. with the teaching of Agrawal et al. "wherein [the] solver is used to solve resource optimization problems." The Examiner further stated that the motivation for combining the references is that "this solver problem solving process provides [a] problem modeler that requires less user skill."

The Applicant respectfully disagrees with the Examiner's reasoning and with the Examiner's conclusion. From the discussion above, it is clear that Agrawal et al. is directed toward a system for classification of data records for data mining purposes. In marked contrast, Clark et al. is directed toward problem modeling. These are entirely different fields of endeavor. As such, there can be no motivation to combine these references. Further, Clark et al. discloses both a problem modeler and a problem solver and states that its system requires less user skill. According to the Examiner's stated motivation of achieving a problem solver that requires less user skill, there would be no motivation to modify Clark et al. at all since Clark et al. by itself purportedly already achieves the Examiner's stated motivation. Rather, the Examiner is clearly using the Applicant's disclosure with the benefit of hindsight in an impermissible attempt to reconstruct the Applicant's invention from the cited references.

For at least the reasons given above, the Applicant submits that claims 1, 6-11 and 35 are allowable over the Agrawal et al. and Clark et al. references.

"In order to rely on a reference as a basis for rejection of an applicant's invention, the reference must either be in field of applicant's endeavor or, if not, then be reasonably pertinent to the particular problem with which the inventor was concerned." Manual of Patent Examining Procedure, Section 2141.01(a) (Rev. 2, May 2004). The Applicant's invention is directed toward methods and apparatus for data store assignment for a data storage system design. See, for example, Applicant's specification at page 3, lines 27-28; page 4, lines 10-13; page 7, lines 23-29; and claims 1, 13, 21 and 30. As discussed above, the Agrawal et al. reference is directed toward a system for classification of data records for data mining purposes. These are

entirely different fields of endeavor; the Agrawal et al. reference is simply not reasonably pertinent to the particular problem to which the applicant was concerned - data storage system design. This is another reason why the Agrawal et al. and Clark et al. references cannot be properly combined and is, therefore, another reason why claims 1, 6-11 and 35 are allowable over them.

Further, even if the Agrawal et al. and Clark et al. references could be combined, such a combination would not teach or suggest all of the limitations of claims 1, 6-11 and 35.

As mentioned, the Applicant's invention is directed toward methods and apparatus for data store assignment for a data storage system design. See, for example, Applicant's specification at page 3, lines 27-28; page 4, lines 10-13; page 7, lines 23-29; and claims 1, 13, 21 and 30. A data "store" may be data storage space, such as a logically contiguous block of storage, striped data storage space or concatenated data storage, that is presented to an application. Applicant's specification at page 3, lines 28-31. A store may be implemented as a logical volume into which data elements, such as files or a database, may be inserted, removed or manipulated. Applicant's specification at page 3, lines 31-33. Each store has requirements, such as capacity and bandwidth requirements. Applicant's specification at page 2, line 33 and page 4, lines 3-9. Storage devices to be included in the design for the storage system are represented by a data structure having nodes. Applicant's specification at page 4, lines 10-17, and Figure 1. Associated with some or all of the nodes of the data structure are attributes which represent various properties, such as the price of a particular device, the capacity of an associated LU or the utilization of SCSI buses attached to a device. Applicant's specification at page 6, lines 20-25. A solver assigns the data stores into the data structure, checking at each node that none of the attributes are exceeded by the requirements of the store. Applicant's specification at page 2, line 34 to page 3, line and at page 6, line 33 to page 7, line 26.

Thus, Applicant's claim 1 recites an apparatus for data store assignment for a data storage system design, comprising: a data structure stored in computer-readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device having respective attributes; representations of a plurality of data stores, each data store having respective requirements; and a solver process stored in computer-readable memory for assigning

the representations of data stores to the data structure based on comparisons of the attributes of the nodes to the requirements of the data stores wherein the solver process makes a first assignment of a store and determines a first metric representative of how well the first assignment meets one or more goals for the data storage system and wherein the solver makes a second assignment of the store and determines a second metric representative of how well the second assignment meets the one or more goals and wherein the solver selects one of the first assignment and the second assignment based on the first and second metrics.

Regarding claim 1, the Examiner stated that "Agrawal et al. teaches an apparatus for data store assignment for a data storage system design...". As explained above, however, Agrawal et al. discloses a decision tree classifier for classifying data records for data mining purposes. Agrawal et al. does not suggest or disclose anything about data store assignment or data storage system design. The Examiner also stated that Agrawal et al. in Figure 3 discloses "a data structure stored in computer readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device having respective attributes...". Agrawal et al. does show a decision-tree in Figure 3, however, this decision tree classifier does not include any nodes that represent "a physical data storage device having respective attributes" as is recited by claim 1. While Agrawal et al. mentions attributes, these are attributes of data records which are used for partitioning the records into subsets; the attributes of Agrawal et al. are not equivalent to the attributes of claim 1 which include attributes of a physical data storage device. The Examiner further stated that Agrawal et al. discloses at col. 1, lines 40-65, "representations of a plurality of data stores, each data store having respective requirements." However, this portion of Agrawal et al. discusses the decision tree classifier. Nowhere does Agrawal et al. suggest or disclose "representations of data stores, each data store having respective requirements," as is recited by claim 1. In sum, Agrawal et al. does not disclose any of these features of claim 1.

The Examiner further stated that Clark et al. at col. 8, lines 54-67 and col. 9, lines 31-64 teaches "a solver-process stored in computer-readable memory for assigning the representations of data stores to the data structure based on comparisons of the attributes of the nodes to the requirements of the data stores wherein the solver process makes a first assignment of a [first] store and determines a first metric representative of how well the first assignment meets one or more goals for the data

storage system and wherein the solver makes a second assignment of the first store and determines a second metric representative of how well the second assignment meets the one or more goals and wherein the solver selects one of the first assignment and the second assignment based on the first and second metrics." However, at col. 8, lines 54-67, Clark et al. discusses a specific embodiment of its problem modeler for telecommunications network design and assign in which the modeler extracts information relevant to this specific problem from a network inventory/topology database. At col. 9, lines 31-64 Clark et al. further discusses the specific embodiment of its problem modeler for telecommunications network design and assign in which a metric generator extracts metrics, such as a "hop count" from a database and populates the solver database in a format expected by the problem solver. Neither of these portions of Clark et al. suggests or discloses the recited solver process of claim 1. In sum, Clark et al. does not suggest or disclose any of these features of claim 1.

In view of the above, it can be seen that Agrawal et al. and Clark et al. are entirely distinct from the present invention as recited by claim 1 and, thus, they clearly do not suggest or disclose the limitations of claim 1. This is another reason why claim 1 is allowable. Claims 6-11 and 35 are allowable at least because they are dependent from an allowable base claim 1.

Claims 13-15, 17, 39, 40, 42 and 43 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Agrawal et al. in view of U.S. Patent No. 6,473,884 issued to Ganai et al. (hereinafter "Ganai et al.").

The Applicant respectfully traverses the rejection on the grounds that Agrawal et al. and Ganai et al. cannot be properly combined. The Applicant also traverses on the grounds that, even if these references could be properly combined, such a combination does not suggest or disclose the Applicant's claimed invention.

Ganai et al. is directed toward a system for equivalence checking of combinatorial circuits. Title of Ganai et al. A first circuit and a second circuit are coupled by an XOR gate to produce an output that represents an error function. Ganai et al., col. 2, lines 55-57. If the output of the XOR gate is ever a logic "1" for any logic combination of inputs, the first circuit and the second circuit are not equivalent. Ganai et al., at col. 2, lines 57-59. This combined circuit allows use of Binary Decision Diagram (BBD) analysis. Ganai et al., col. 2, lines 59-62. Before BDD analysis is performed, the circuit is represented as a combination of logic AND gates and inverters. Ganai et al., at col. 2, lines 62-64. Constructing a BDD model involves

starting with a first input, for each possible state of the circuit inputs, a decision tree is built that terminates in either another decision based on another input or terminates in a determined logic state. Ganai et al., at col. 3, lines 8-20, and Figures 2A-B. A satisfiability (SAT) solver can be used to verify the equivalence of the first and second circuits. Ganai et al., at col. 3, lines 28-32. If the SAT solver is able to find an assignment of values to all nodes in the network such that the XOR gate output is a logic zero, the equivalence of the two circuits is established. Ganai et al., at col. 3, lines 35-38. Operation of the SAT solver is shown in Figure 4 of Ganai et al. and described at col. 4, lines 17-46.

The Examiner stated that it would have been obvious to combine Ganai et al. with Agrawal et al. and provided essentially the same motivation for combining Ganai et al. with Agrawal et al.: "that this solver problem solving process provides [a] problem modeler that requires less user skill." However, neither of these references teaches or suggests less user skill; rather, Clark et al. discusses this feature. Because less user skill is not suggested by either of the Ganai et al. or Agrawal et al. references, this cannot provide a motivation to combine the references. Further, as explained above, Agrawal et al. discloses a decision tree classifier for data mining purposes. Ganai et al. discloses a method and system for equivalence checking of combinatorial circuits. These are entirely different fields of endeavor. As such, there can be no motivation to combine these references.

For at least the reasons given above, the Applicant's claims 13-15, 17, 39, 40, 42 and 43 are allowable over the Agrawal et al. and Ganai et al. references.

Further, as explained above, the Applicant's invention is directed toward methods and apparatus for data store assignment for a data storage system design. And, as explained above, the Agrawal et al. reference is directed toward a system for classification of data records for data mining purposes. The Ganai et al. reference is directed toward equivalence checking of combinatorial circuits. These are entirely different fields of endeavor; neither the Agrawal et al. reference, nor the Ganai et al. reference reasonably pertinent to the particular problem to which the applicant was concerned - data storage system design. This is another reason why the Agrawal et al. and Ganai et al. references cannot be properly combined and is, therefore, another reason why claims 13-15, 17, 39, 40, 42 and 43 are allowable.

Further, even if the Agrawal et al. and Ganai et al. references could be combined, such a combination would not teach or suggest all of the limitations of the Applicant's claims 13-15, 17, 39, 40, 42 and 43.

The Applicant's claim 13 recites an apparatus for data store assignment for a data storage system design, comprising: a data structure stored in computer-readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device and having respective attributes; representations of a plurality of data stores, each data store having respective requirements; and a solver process stored in computer-readable memory for assigning the representations of data stores to the nodes wherein the solver process compares the requirements of a store to the attributes of one or more of the data storage devices and when the comparison indicates compatibility, the solver process makes a first assignment of the store and when the comparison does not indicate compatibility, the solver process modifies one or more of the attributes of one or more of the data storage devices.

Regarding claim 13, the Examiner stated that "Agrawal et al. teaches an apparatus for data store assignment for a data storage system design...". As explained above, however, Agrawal et al. discloses a decision tree classifier for classifying data records for data mining purposes. Agrawal et al. does not suggest or disclose anything about data store assignment or data storage system design. The Examiner also stated that Agrawal et al. in Figure 3 discloses "a data structure stored in computer readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device having respective attributes...". Agrawal et al. does show a decision-tree in Figure 3, however, this decision tree classifier does not include any nodes that represent "a physical data storage device having respective attributes" as is recited by claim 13. While Agrawal et al. mentions attributes, these are attributes of data records which are used for partitioning the records into subsets; the attributes of Agrawal et al. are not equivalent to the attributes of claim 13 which include attributes of a physical data storage device. The Examiner further stated that Agrawal et al. discloses at col. 1, lines 40-65, "representations of a plurality of data stores, each data store having respective requirements." However, this portion of Agrawal et al. discusses the decision tree classifier. Nowhere does Agrawal et al. suggest or disclose "representations of data

stores, each data store having respective requirements," as is recited by claim 13. In sum, Agrawal et al. does not disclose any of these features of claim 13.

The Examiner further stated that at col. 3, lines 22-43 and col. 4, lines 17-44 Ganai et al. teaches "a solver process stored in computer-readable memory for assigning the representations of data stores to the nodes wherein the solver process compares the requirements of a store to the attributes of one or more of the data storage devices and when the comparison indicates compatibility, the solver process makes a first assignment of the store and when the comparison does not indicate compatibility, the solver process modifies one or more of the attributes of one or more of the data storage devices." However, at col. 3, lines 22-43 Ganai et al. discusses use of its BDD analysis and SAT solver to determine whether two combinatorial circuits are equivalent. At col. 4, lines 17-44 Ganai et al. discusses operation of the SAT solver which is used to verify the equivalence of the first and second combinatorial circuits. Neither of these portions of Ganai suggests or discloses the recited solver process of claim 13. In sum, Clark et al. does not suggest or disclose any of these features of claim 13.

In view of the above, it can be seen that Agrawal et al. and Clark et al. are entirely distinct from the present invention as recited by claim 13 and, thus, they clearly do not suggest or disclose the limitations of claim 13. This is yet another reason why claim 13 is allowable. Claims 14-15, 17, 39 and 40 are allowable at least because they are dependent from an allowable base claim 13.

The Applicant's claim 42 recites an apparatus for data store assignment for a data storage system design, comprising: a data structure stored in computer-readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device and having respective attributes; representations of a plurality of data stores, each data store having respective requirements; and a solver process stored in computer-readable memory for assigning the representations of data stores to the nodes wherein the solver process compares the requirements of a store to the attributes of one or more of the data storage devices and when the comparison indicates compatibility, the solver process makes a first assignment of the store and when the comparison does not indicate compatibility, the solver process adds one or more additional nodes to the data structure.

Regarding claim 42, the Examiner stated that "Agrawal et al. teaches an apparatus for data store assignment for a data storage system design...". As explained

above, however, Agrawal et al. discloses a decision tree classifier for classifying data records for data mining purposes. Agrawal et al. does not suggest or disclose anything about data store assignment or data storage system design. The Examiner also stated that Agrawal et al. in Figure 3 discloses "a data structure stored in computer readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device having respective attributes...". Agrawal et al. does show a decision-tree in Figure 3, however, this decision tree classifier does not include any nodes that represent "a physical data storage device having respective attributes" as is recited by claim 42. While Agrawal et al. mentions attributes, these are attributes of data records which are used for partitioning the records into subsets; the attributes of Agrawal et al. are not equivalent to the attributes of claim 42 which include attributes of a physical data storage device. The Examiner further stated that Agrawal et al. discloses at col. 1, lines 40-65, "representations of a plurality of data stores, each data store having respective requirements." However, this portion of Agrawal et al. discusses the decision tree classifier. Nowhere does Agrawal et al. suggest or disclose "representations of data stores, each data store having respective requirements," as is recited by claim 42. In sum, Agrawal et al. does not suggest or disclose any of these features of claim 42.

The Examiner further stated that at col. 3, lines 22-43 and col. 4, lines 17-44 Ganai et al. teaches "a solver process stored in computer-readable memory for assigning the representations of data stores to the nodes wherein the solver process compares the requirements of a store to the attributes of one or more of the data storage devices and when the comparison indicates compatibility, the solver process makes a first assignment of the store and when the comparison does not indicate compatibility, the solver process adds one or more additional nodes to the data structure." However, as explained above, at col. 3, lines 22-43 Ganai et al. discusses use of its BDD analysis and SAT solver to determine whether two combinatorial circuits are equivalent. At col. 4, lines 17-44 Ganai et al. discusses operation of the SAT solver which is used to verify the equivalence of the first and second combinatorial circuits. Neither of these portions of Ganai suggests or discloses the recited solver process of claim 42. In sum, Clark et al. does not suggest or disclose any of these features of claim 42.

In view of the above, it can be seen that Agrawal et al. and Clark et al. are entirely distinct from the present invention as recited by claim 42 and, thus, they

clearly do not suggest or disclose the limitations of claim 42. This is yet another reason why claim 42 is allowable. Claims 43 is allowable at least because it is dependent from an allowable base claim 42.

Claims 21, 23-25, 30, 32 and 33 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Ganai et al. in view of Clark et al.

The Applicant respectfully traverses the rejection on the grounds that Ganai et al. and Clark et al cannot be properly combined. The Applicant also traverses on the grounds that, even if these references could be properly combined, such a combination does not suggest or disclose the Applicant's claimed invention.

As explained above, Ganai et al. discloses a method and system for equivalence checking of combinatorial circuits. As is also explained above, Clark et al. is directed toward problem modeling. These are entirely different fields of endeavor. As such, there can be no motivation to combine these references.

The Examiner stated that it would have been obvious to combine Ganai et al. with Clark et al. "wherein problem solver determines metrics and metric relationship" and asserted that the motivation for combining Ganai et al. with Clark et al. is that "metric generator periodically or continuously monitors information resource to determine its current objects and the state of those objects, and then updates the metrics for solver database." The feature of a metric generator is disclosed by Clark et al. and is used in the system of Clark et al. as part of its problem modeler to extract information from the resource information database for use by the problem solver of Clark et al. See col. 9, lines 31-64 of Clark et al. As explained above, Ganai et al. uses Binary Decision Diagrams (BDDs) and a satisfiability (SAT) solver to determine whether two combinational circuits are equivalent. Ganai et al. is not directed to problem modeling and, thus, does not employ a problem modeler. Accordingly, there would be no purpose for using such a metric generator in Ganai et al. As such, this cannot be a motivation to combine these references.

In view of the above, the Clark et al. and Ganai et al references cannot be properly combined. Accordingly, claims 21, 23-25, 30, 32 and 33 are allowable.

Further, even if the Ganai et al. and Clark et al. references could be combined, such a combination would not teach or suggest all of the limitations of the Applicant's claims 21, 23-25, 30, 32 and 33.

Regarding claim 21, the Examiner stated that "Ganai teaches a method of data store assignment for a data storage system...". As explained above, however, Ganai

et al. discloses a method and system for equivalence checking of combinatorial circuits. Ganai et al. do not suggest or disclose anything about data store assignment or data storage system design. The Examiner also stated that Ganai et al. in col. 3, lines 3-32 discloses "providing a data structure stored in computer readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device having respective attributes...". This portion of Ganai et al. does discuss a decision-tree in connection with its Binary Decision Diagram (BDD) model of a combinatorial circuit; however, this BDD model does not include any nodes that represent "a physical data storage device having respective attributes" as is recited by claim 21. The Examiner further stated that Ganai et al. discloses at col. 3, lines 32-35, "providing a representation of a first data store, having requirements." However, this portion of Ganai et al. discusses the satisfiability SAT solver which determines equivalence of the first and second circuits of Ganai et al. This simply does not suggest or disclose the recited feature of claim 21.

The Examiner also stated that Ganai et al. discloses at col. 3, lines 3-6, "comparing the requirements for the data store to attributes of devices in the data structure." However, this portion of Ganai et al. discloses its verification analysis in which combinatorial circuits are said to be usually identical with the exception of errors. This has nothing to do with comparing data store requirements with device attributes, as recited by claim 21. The Examiner further stated that Ganai et al. discloses at col. 3, lines 22-43 "making a first assignment of the representation of the first data store to the data structure based on results of said comparing." However, this portion of Ganai et al. discusses use of its BDD analysis and SAT solver to determine whether the two combinatorial circuits are equivalent. This simply does not suggest or disclose the recited feature of claim 21. The Examiner further stated that Ganai et al. discloses at col. 4, lines 7-16 "making a second assignment of the representation of the data store based on said results of said comparing." However, this portion of Ganai et al. further discusses use of its BDD analysis and SAT solver to determine whether the two combinatorial circuits are equivalent. Again, this simply does not suggest or disclose the recited feature of the claim 21. In sum, Ganai et al. does not suggest or disclose any of these features of claim 21.

The Examiner further stated that Clark et al. at col. 8, lines 54-67 teaches "determining a first metric for the first assignment meets one or more goals for the

data storage system." However, this portion of Clark et al. discusses a specific embodiment of its problem modeler for telecommunications network design and assign in which the modeler extracts information relevant to this specific problem from a network inventory/topology database. This does not suggest or disclose the recited feature of claim 21. The Examiner further stated that Clark et al. at col. 9, lines 31-64 discloses "determining a second metric for the second assignment according to how well the second assignment meets one or more goals for the data storage system." However, this portion of Clark et al. further discusses the specific embodiment of its problem modeler for telecommunications network design and assign in which a metric generator extracts metrics, such as a "hop count" from a database and populates the solver database in a format expected by the problem solver. This simply does not suggest or disclose the recited feature of claim 21. The Examiner further stated that Clark et al. at col. 9, lines 51-64 discloses "selecting the first assignment or the second assignment based on the first and second metrics." However, this portion of Clark et al. further discusses the specific embodiment of its problem modeler for telecommunications network design and assign in which a metric generator monitors the network topology database to determine its current objects and the state of those objects and updates the solver database so that it has the most current transformed view of the topology database. Again, this does not suggest or disclose the recited feature of claim 21. In sum, Clark et al. does not suggest or disclose any of these features of claim 21.

In view of the above, it can be seen that Ganai et al. and Clark et al. are entirely distinct from the present invention as recited by claim 21 and, thus, they clearly do not suggest or disclose the limitations of claim 21. This is another reason why claim 21 is allowable. Claims 23-25 are allowable at least because they are dependent from an allowable base claim 21.

Regarding claim 30, the Examiner stated that "Ganai teaches a method of data store assignment for a data storage system...". As explained above, however, Ganai et al. discloses a method and system for equivalence checking of combinatorial circuits. Ganai et al. clearly do not suggest or disclose anything about data store assignment or data storage system design. The Examiner also stated that Ganai et al. in col. 3, lines 3-32 discloses "providing a data structure stored in computer readable memory, the data structure having a plurality of nodes, at least some of the nodes each representing a physical data storage device having respective attributes...". As

explained above, this portion of Ganai et al. does discuss a decision-tree in connection with its Binary Decision Diagram (BDD) model of a combinatorial circuit; however, this BDD model does not include any nodes that represent "a physical data storage device having respective attributes" as is recited by claim 30. The Examiner further stated that Ganai et al. discloses at col. 3, lines 32-35, "providing a representations of a plurality of data stores, each having requirements." However, this portion of Ganai et al. discusses the satisfiability SAT solver which determines equivalence of the first and second circuits of Ganai et al.

The Examiner also stated that Ganai et al. discloses at col. 3, lines 3-6, "comparing the requirements for each data store to attributes of devices in the data structure." As explained above, this portion of Ganai et al. discloses its verification analysis in which combinatorial circuits are said to be usually identical with the exception of errors. This simply has nothing to do with comparing data store requirements with device attributes, as recited by claim 30. The Examiner further stated that Ganai et al. discloses at col. 3, lines 22-43, "making a first assignment of the representations of each of the plurality of stores to the data structure based on results of said comparing." As explained above, this portion of Ganai et al. discusses use of its BDD analysis and SAT solver to determine whether the two combinatorial circuits are equivalent. This does not suggest or disclose the recited feature of claim 30. The Examiner further stated that Ganai et al. discloses at col. 4, lines 7-16 "making a second assignment of the representations of each of the plurality of data stores based on said results of said comparing." As explained above, this portion of Ganai et al. further discusses use of its BDD analysis and SAT solver to determine whether the two combinatorial circuits are equivalent. Again, this does not suggest or disclose the recited feature of claim 30. In sum, Ganai et al. does not suggest or disclose any of these features of claim 30.

The Examiner further stated that Clark et al. at col. 8, lines 54-67, teaches "determining a first metric for the first assignment meets one or more goals for the data storage system." As explained above, this portion of Clark et al. discusses a specific embodiment of its problem modeler for telecommunications network design and assign in which the modeler extracts information relevant to this specific problem from a network inventory/topology database. This does not suggest or disclose the recited feature of claim 30. The Examiner further stated that Clark et al. at col. 9, lines 31-64 discloses "determining a second metric for the second assignment

according to how well the second assignment meets one or more goals for the data storage system." As explained above, this portion of Clark et al. further discusses the specific embodiment of its problem modeler for telecommunications network design and assign in which a metric generator extracts metrics, such as a "hop count," from a database and populates the solver database in a format expected by the problem solver. This does not suggest or disclose the recited feature of claim 30. The Examiner further stated that Clark et al. at col. 9, lines 51-64 discloses "selecting the first assignment or the second assignment based on the first and second metrics." As explained above, this portion of Clark et al. further discusses the specific embodiment of its problem modeler for telecommunications network design and assign in which a metric generator monitors the network topology database to determine its current objects and the state of those objects and updates the solver database so that it has the most current transformed view of the topology database. Again, this does not suggest or disclose the recited feature of claim 30. In sum, Clark et al. does not suggest or disclose any of these features of claim 30.

In view of the above, it can be seen that Ganai et al. and Clark et al. are entirely distinct from the present invention as recited by claim 30 and, thus, they clearly do not suggest or disclose the limitations of claim 30. This is another reason why claim 30 is allowable. Claims 32 and 33 are allowable at least because they are dependent from an allowable base claim 30.

The Examiner also rejected dependent claims 2-5, 12, 16, 18-20, 22, 31, 34 36-38 and 41 in view of Ganai, Clark and Agrawal, all taken in combination. These claims are all dependent from allowable base claims. For at least this reason, these claims are allowable.

Conclusion:

The Agrawal et al. reference discloses construction of a decision tree classifier by recursively partitioning a set of training records. The classifier is used for classifying data records for data mining purposes. The Clark et al. reference discloses a generic problem modeler for examining an existing user information resource and for transforming the relevant information from the resource into data that can be stored in a solver database which is directly accessible by a problem solver. The Ganai et al. reference discloses a system for equivalence checking of combinatorial

circuits using Binary Decision Diagram (BBD) analysis and a satisfiability (SAT) solver. Theses references are from entirely different fields of endeavor and, thus, they cannot be properly combined. Moreover, they do not suggest or disclose the Applicant's claimed invention which is directed toward data store assignment for data storage system design. Rather, the Examiner is using the Applicant's disclosure with the benefit of hindsight in an impermissible attempt to reconstruct the Applicant's invention from the cited references.

In view of the above, the Applicant submits that all of the pending claims are now allowable. Allowance at an early date would be greatly appreciated. Should any outstanding issues remain, the examiner is encouraged to contact the undersigned at (408) 293-9000 so that any such issues can be expeditiously resolved.

Respectfully Submitted,

Dated: Sept 3, 2004

Derek J. Westberg (Reg. No. 40,872)